ON STOCHASTIC CONTROL IN FINANCE

WOLFGANG J. RUNGGALDIER*

Abstract. Stochastic control/optimization problems arise in various applications in finance where the control is usually given by an investment strategy. The purpose of this paper is to review some of these applications together with appropriate solution methodologies and also to discuss the latter in comparison with one another.

Key words. Financial markets, optimal investment strategies, hedging, stochastic control, model uncertainty, incomplete information, adaptive approaches, robust approaches.

AMS(MOS) subject classifications. Primary 91B28, 91B70, 93E20; Secondary 93E35, 93E11, 91B30, 90B50.

1. Introduction.

The purpose of this paper is to review some financial applications of stochastic control together with appropriate solution methodologies as they have appeared in the recent literature. No claim is made of completeness and the treatment may partly reflect the specific interests and competencies of the author. Since for a given problem there may be various competing methodologies, wherever possible we shall also discuss them in comparison with one another.

On the side of the reader we assume some basic knowledge of stochastic control and stochastic calculus as well as of Dynamic Programming methodologies. Concerning the financial applications we have however tried to make the paper as self contained as possible. For this purpose we have devoted a fair amount of space, more precisely the sections from 1 through 4, to a description of the financial models and problems and to the discussion of the relevant solution approaches referring for more details to the existing literature. In Sections 5 and 6 we then present results more specific to the author's own interests. In particular, in Section 5 we treat cases of incomplete information (*model risk* in the financial terminology) while in Section 6 we briefly describe two additional topics, one concerning almost sure optimality and the other an application of stochastic filtering.

2. Models and problems.

Optimization problems in finance are mainly related to investment problems. In fact, agents invest in the various assets available on the market with the objective of achieving specific goals. Before mentioning such possible goals, we recall some basic market models, where one typically considers a locally riskless asset and a certain number of risky assets. The evolution of the corresponding prices may be described either in continuous or discrete time. We start by recalling continuous time

*Dipartimento di Matematica Pura ed Applicata, Universitá di Padova, Via Belzoni 7, 35131-Padova, Italy (runggal@math.unipd.it; http://www.math.unipd.it/runggaldier/index.html).

models, which is the traditional setup. However, since investment decisions take place in discrete time, it makes sense to consider also discrete time models that we shall briefly recall as well.

2.1. Financial market models (continuous time). As mentioned, there is always a locally riskless asset that may be viewed as a bank account/money market and that is supposed to evolve according to

$$dB_t = r_t B_t dt. \tag{2.1}$$

If the rate r_t (*spot interest rate*) is a stochastic process, also B_t is risky. Locally, at time t, it is however riskless. By the economic principle of absence of arbitrage, the growth rate of any locally riskless asset has to be unique and so it suffices to consider only one such asset.

For the risky assets, for which the prices will be denoted by $S_t = [S_t^1, \cdots, S_t^N]'$ for some $N \in \mathbb{N}$, the basic evolution model is that of a lognormal diffusion process

$$dS_t^i = S_t^i a_t^i dt + S_t^i \sum_{j=1}^{M} \sigma_t^{i,j} dw_t^j \tag{2.2}$$

where, for some $M \in \mathbb{N}$, $w_t = [w_t^1, \cdots, w_t^M]'$ is a Wiener process on a filtered probability space $(\Omega, \mathcal{F}, \mathcal{F}_t, P)$ and a_t^i as well as $\sigma_t^{i,j}$ are deterministic. In vector notation (2.2) may be represented as

$$dS_t = (diag\, S_t)\, A_t dt + (diag\, S_t)\, \Sigma_t dw_t \tag{2.3}$$

where by $(diag\, S)$ we denote the $N \times N-$diagonal matrix with diagonal elements $S_t^i (i = 1, \cdots, N)$. More generally, the coefficients may be stochastic processes, adapted to $\mathcal{F}_t$ or to a larger filtration (see point a) below), in which case S_t^i are no longer lognormal, but they take always positive values as it should be since they describe prices.

Agents invest according to an investment strategy that can be described in two equivalent ways. The first is to consider a strategy as an $(N+1)-$dimensional, $\mathcal{F}_t-$predictable process

$$\bar{\phi}_t = (\phi_t^0, \phi_t^1, \cdots, \phi_t^N) := (\phi_t^0, \phi_t) \tag{2.4}$$

where ϕ_t^i denotes the n*umber of shares* of asset i held in the portfolio at time t ($i = 0$ refers to the riskless asset). The value of the corresponding portfolio is then given by

$$V_t = \phi_t^0 B_t + \sum_{i=1}^{N} \phi_t^i S_t^i. \tag{2.5}$$

Equivalently, one may consider the vector

$$\pi_t = (\pi_t^1, \cdots, \pi_t^N) \tag{2.6}$$

where $\pi_t^i = \frac{\phi_t^i S_t^i}{V_t}$, $(i = 1, \cdots, N)$ denotes the *fraction of wealth* invested in the risky asset i at time t. Given π_t, the fraction of wealth invested in the riskless asset at time t is then $1 - \sum_{i=1}^N \pi_t^i$. Depending on the context, we shall occasionally denote the portfolio value also by V_t^ϕ or V_t^π.

An important class of strategies is that of the *self financing strategies*, which can be expressed more intuitively in terms of ϕ_t as

$$dV_t = \phi_t^0 dB_t + \sum_{i=1}^N \phi_t^i dS_t^i. \tag{2.7}$$

Relation (2.7) expresses the fact that, for such strategies, the infinitesimal change over time in the portfolio value is due only to changes in the prices of the assets in the portfolio. In other words, the portfolio may be rebalanced over time, but without adding funds to it nor consuming from it. Notice that $\phi_t^0 B_t = V_t - \phi_t S_t$, $(\phi_t S_t = \sum_{i=1}^N \phi_t^i S_t^i)$, which, using (2.1) and (2.3), allows one to express the self financing condition (2.7) in vector notation as

$$\begin{aligned} dV_t &= r_t \left[V_t - \phi_t S_t\right] dt + \phi_t dS_t \\ &= V_t \left[r_t dt + \pi_t (A_t - r_t \underline{1}) dt + \pi_t \Sigma_t dw_t\right]. \end{aligned} \tag{2.8}$$

Notice that, expressing the self financing condition in terms of π_t as in (2.8), allows one to obtain the dynamics of V_t autonomously without reference to the individual assets B_t and S_t.

Often it is convenient to consider the prices discounted with respect to the nonrisky asset, i.e. $\tilde{S}_t^i := B_t^{-1} S_t^i$ (analogously, $\tilde{V}_t := B_t^{-1} V_t$) for which (2.7) simplifies to

$$d\tilde{V}_t = \phi_t d\tilde{S}_t. \tag{2.9}$$

Below we shall occasionally use S_t and V_t to denote also the discounted values when this does not create possible confusion.

Remark 2.1. The self financing condition, expressed in (2.7) or (2.8), does not take into account possible transaction costs. These have been dealt with in the literature only more recently (among the earlier papers are [56] and [21] in a more economic, respectively mathematical context). Their impact is less severe in discrete time than in continuous time; in the latter case an optimal investment strategy may in fact require a continuous rebalancing and, with proportional transaction costs, this implies that the total cost grows as the quadratic variation of the prices of the risky assets, i.e. to infinity. For simplicity, we shall in general not deal with transaction costs here either; for the moment we limit ourselves to mention a possible adaptation to transaction costs of (2.7) in a simplified context when there is no bartering between the risky assets. In other words, we shall assume that an investor must purchase risky assets through money. Assume then that L_t^i denotes the cumulative number of shares of the $i-$th asset purchased up to time t and, analogously, M_t^i for those sold. Assume also that λ^i and μ^i denote the cost rate (with respect to the amount traded) for buying, respectively selling asset i. With discounted costs one can then generalize (2.7) to become (see e.g. [38])

$$dV_t = \phi_t^0 dB_t + \sum_{i=1}^{N} \left[\phi_t^i dS_t^i - S_t^i \left(\lambda^i dL_t^i + \mu^i dM_t^i\right)\right] \tag{2.10}$$

where now $\phi_t^i = L_t^i - M_t^i$. Among the more recent studies concerning transaction costs we just mention [82, 17, 42, 43, 8, 70] and the references therein.

Model (2.2) is the traditional lognormal asset price model, that corresponds to the classical Black and Scholes model [7]. Among the various possible generalizations we mention here just a few of the more significant ones.

a) Markov modulated coefficients (factor models). Corresponding to (2.3) let

$$dS_t = (diag\, S_t)\, A_t(Z_t)dt + (diag\, S_t)\, \Sigma_t(Z_t)dw_t \tag{2.11}$$

where Z_t is a (multivariate) Markov process that is generally adapted to a filtration larger than $\mathcal{F}_t^w$ and that may be observable or not (*hidden Markov model*). The process Z_t may describe what is often called the *exogenous factors*, which may have various economic interpretations such as the level of interest rates or, more generally, as the "state of the economy", but also as e.g. the volume of trade, etc. Model (2.11) corresponds also to what in the more engineering literature is called a *regime switching model*. When Z_t is not directly observable, it can under certain conditions be estimated from the observations of S_t through the quadratic variation of S_t. In general however this will not be possible and so Z has to be filtered from the $S-$observations.

b) Jump-diffusion models. The trajectories generated by (2.2) (equivalently (2.3)) are continuous in time, while in reality prices may exhibit a jumping behavior. This may be taken into account by considering in (2.3) also an additive jump term, which in the simplest case takes the form

$$dS_t = (diag\, S_t)\, A_t dt + (diag\, S_t)\, \Sigma_t dw_t + (diag S_{t-})\Psi_{t-} dN_t \tag{2.12}$$

with $N_t = (N_t^1, \cdots, N_t^K)'$ a $K-$dimensional Poisson jump process and Ψ_t an $(N \times K)-$matrix with elements $\Psi_t^{i,j} > -1$ so that the solution is positive. A first study of such models in a financial context goes back to [60]. The model may be further generalized to driving marked point processes (see e.g. the survey in [67]) as well as to Levy driving processes (see e.g. [5, 44]). Jump-diffusion models may on the other hand be contrasted by models of the type (2.2), where the driving noise is a fractional Brownian motion (see e.g. [57, 26, 40]).

c) Models for high frequency data. On small time scales the prices do not change continuously over time, but rather at discrete random points in time when new significant information is being absorbed by the market and the market makers update their quotes. A possibility to model such a situation is by sampling the values of the price process S_t (given by (2.2) or its generalizations (2.11) and (2.12)) at the jump times of a Poisson process N_t. If, as in (2.11), there is a (hidden) Markov process that affects the coefficients in the $S-$dynamics, then it is reasonable to assume that this Z_t affects also the intensity of the Poisson jump process N_t. For a more detailed description of such a model as well as an example of application of stochastic filtering to such a setup we refer to [35].

We conclude this section by pointing out that the assets in (2.2) do not generally represent all possible assets available on the market, but only those in which an agent intends to invest. One may thus consider, in addition to actual portfolios, theoretical (self financing) portfolios that include all or most of the assets available on a given market. Such portfolios may serve as index or *benchmark* (e.g. the S&P 500 index) and one of the goals of an investor may be to form an actual portfolio, with a given selection of assets, that tracks or "beats" in some sense the benchmark. Denoting by I_t the value in t of such a benchmark portfolio and assuming that its (theoretical) portfolio strategy is expressed in terms of the fraction of wealth invested in the various assets, we have that its dynamics may be described by analogy to (2.8) as

$$dI_t = I_t\left[\alpha_t dt + \sigma_t dw_t + \rho_t dv_t\right] \tag{2.13}$$

with predictable coefficients and where w_t is the (multivariate) Wiener process driving the assets chosen for actual investment, while v_t is a Wiener process independent of w_t and resulting from the disturbances that drive exclusively the additional assets on the market. Although correlated with the investment opportunities, the index I_t is however not perfectly correlated with the latter which implies that, even if the market without the index is complete, with the index it is incomplete so that perfect tracking is not possible. In this context notice also that, even if an investor knows the composition of the benchmark portfolio, the best strategy to track the benchmark is not to invest in the chosen assets with the same proportions as they appear in the benchmark, unless the additional assets are independent of the chosen ones. We also point out that the description of a benchmark as a theoretical portfolio, which includes all/most assets, is not the only possible one. In an insurance context it may e.g. represent a wage process, but it may also represent inflation etc. As a reference on such problems we mention [9].

2.2. Financial market models (discrete time). Since investment decisions take place in discrete time, it suffices to know the prices of the assets at the discrete time points where an investment decision takes place. This justifies, from a modelling point of view, the use also of discrete time models; the latter may, furthermore, be useful as approximations of continuous time models when it comes to actual numerical computations. The general setup is completely analogous to that in continuous time and so we mention here only some more specific aspects, letting Δ denote the discrete time step.

As in continuous time, there is one locally riskless asset, the price of which evolves by analogy to (2.1) according to

$$B_{n+1} = B_n(1 + r_n). \tag{2.14}$$

Concerning the risky assets, from (2.2) one has

$$\begin{aligned} S^i_{n+1} = S^i_n \exp\left\{ \int_{n\Delta}^{(n+1)\Delta} \left(a^i_s - \frac{1}{2}\sum_{j=1}^{M} (\sigma^{i,j}_s)^2 \right) ds \right. \\ \left. + \sum_{j=1}^{M} \int_{n\Delta}^{(n+1)\Delta} \sigma^{i,j}_s dw^j_s \right\} \\ := S^i_n \xi^i_n \end{aligned} \tag{2.15}$$

where (ξ^i_n) is thus, for each $i = 1, \cdots, N$, a sequence of lognormal random variables. Relation (2.15) can now be taken as the discrete time evolution model for the prices of the risky assets, that is completely analogous to (2.14) with the deterministic sequence $(1 + r_n)$ replaced by the stochastic sequence ξ^i_n. Notice that, if (2.15) is derived from the more general jump-diffusion model (2.12), then the corresponding discrete time model remains of the same form $S^i_{n+1} = S^i_n \xi^i_n$, except that the ξ^i_n are no longer lognormal. If one starts instead from the Markov modulated model (2.11), then (see e.g. [69]) the ξ^i_n are lognormal conditionally on Z_n and with mean and variance that is a function of Z_n. Typically, one would further discretize (in space) the random variables ξ^i_n, thus ending up with a multinomial model, of which the classical binomial or CRR model (see [13]) is a particular case.

Concerning investment strategies, recall that in continuous time we had assumed them to be predictable. In discrete time this implies that the strategy for period $n + 1$ has to be determined in period n, whereby we consider as "period n" the interval $[(n - 1)\Delta, n\Delta)$. Assuming that prices are already discounted with respect to B_n, the self financing condition corresponding to the more general setup in (2.10), which includes proportional transaction costs, can then be expressed as follows, where ϕ^i_n denotes the number of shares of asset i held in the portfolio in period n

$$V_{n+1} = V_n + \sum_{i=1}^{N} S^i_n \left\{ \phi^i_{n+1}(\xi^i_n - 1) - \lambda^i(\phi^i_{n+1} - \phi^i_n)^+ - \mu^i(\phi^i_{n+1} - \phi^i_n)^- \right\} \tag{2.16}$$

A more intuitive justification of (2.16) can be seen in the following equivalent relation for the case when there are no transaction costs

$$\phi^0_{n+1} + \sum_{i=1}^{N} \phi^i_{n+1} S^i_n = \phi^0_n + \sum_{i=1}^{N} \phi^i_n S^i_n \tag{2.17}$$

which helps also to clarify the use of ϕ^i_{n+1} (not ϕ^i_n) in the second term on the right hand side in (2.16).

3. Control/optimization problems. In this section we mention some of the basic control/optimization problems in finance, all investment-type problems, thereby concentrating ourselves mainly on a particular class, namely hedging problems.

In any investment problem the control is given by the investment strategy. We shall consider only *small investors*, whose investment decisions do not affect the prices on the market. As a consequence, one may consider two groups of *state variables* : on one hand the asset prices and a possible benchmark process that are uncontrolled, on the other hand the portfolio value process that is the only controlled state variable. The objective function depends usually on both types of state variables.

3.1. Basic control/optimization problems.

A) The classical optimization problem is that of **maximizing utility from consumption and terminal wealth**. We recall here its standard formulation in continuous time that goes back to [59], with S_t evolving according to (2.3) and V_t according to (2.8), thereby neglecting transaction costs but considering an additional control variable c_t that represents the rate of consumption at time t:

$$\begin{cases} dS_t = (diag\, S_t)\, A_t dt + (diag\, S_t)\, \Sigma_t dw_t \\ dV^\pi_t = V^\pi_t \left[r_t dt + \pi_t (A_t - r_t \underline{1}) dt + \pi_t \Sigma_t dw_t \right] - c_t dt \\ \max\limits_{\pi} E_{S_0, V_0} \left\{ \int_0^T U_1(t, c_t) dt + U_2(V^\pi_T) \right\}. \end{cases} \tag{3.1}$$

Here $U_1(\cdot)$ and $U_2(\cdot)$ are utility functions from consumption and terminal wealth respectively and they satisfy the usual assumptions. Relation (3.1) defines now a standard finite horizon stochastic control problem with a running cost and a terminal cost. Of course, one may consider all the standard variants of this basic problem.

B) An area where stochastic control/optimization problems arise quite naturally is **dynamic risk management for an insurance company**. The basic goal is that of *minimizing ruin probability/maximizing survival probability*, but also other goals can be considered such as maximizing dividend payout. Typical control variables are the amount of reinsurance, proportional and not, as well as investment for the case when an insurer has the possibility of investing part of his surplus into a risky asset. We shall not enter into details here, but refer to the more specific literature such as e.g [58, 62, 39, 3, 12, 74, 37] and the references therein.

C) A problem class that has attracted recently considerable interest is that of **hedging a future liability** or the related problem of **tracking a benchmark** and this will be the subject of the next subsection.

3.2. The problem of hedging a future liability/contingent claim. We shall describe this problem in a continuous time setup with S_t evolving according to (2.3) and V_t according to (2.8). We shall also discuss various issues in the context of this problem, but they may apply also to the other problem-classes.

Suppose that at a given future date T (*maturity*) we face a random liability, modelled as a random variable H_T adapted to $\mathcal{F}_T$. At T the amount of this liability becomes thus known, but for $t < T$ it is a random variable that depends on the evolution of the random sources that drive the market (*contingent claim*). The liability implies thus some risk and one purpose of an investor is to hedge this risk with a portfolio resulting from investing in the assets on the market. As in the rest of the paper we shall neglect transaction costs; in the present context their presence not only complicates the dynamics in (2.8), but requires also to specify whether the liability has to be settled in cash or otherwise.

The ideal situation is that of being able to determine an initial capital and a self financing strategy π such that, at the maturity T,

$$V_T^\pi = H_T \ , \quad \text{a.s.} \tag{3.2}$$

This is the case of **perfect hedging/replication/duplication**, which is possible only in so-called complete markets that represent an idealization of reality.

If perfect hedging is not possible, one can use various alternative criteria such as

A) Superhedging. Determine a self financing π and the smallest initial capital $V_0 = V_0^\pi$ such that

$$V_T^\pi \geq H_T \quad \text{a.s.} \tag{3.3}$$

The a.s. requirement means that the investor wants to protect himself also in extreme events and this requires a large initial capital. Among the various studies related to this criterion we mention here just [23, 2, 27, 52] as well as [79] that concerns also an alternative to dealing with transaction costs, namely the bounding of the variation of the optimal strategy (*gamma constraint*).

An alternative criterion is to minimize a downside-type risk, namely

B) Minimization of shortfall risk. For given S_0 and V_0 determine a self financing ϕ such that

$$E_{S_0,V_0}\left\{\mathcal{L}\left[(H_T - V_T^\pi)^+\right]\right\} \quad \to \quad \min \tag{3.4}$$

where $\mathcal{L}(\cdot)$ is a given loss function and $(\cdot)^+$ denotes the positive part. The quantity $(H_T - V_T^\pi)^+$ is in fact the amount by which the investor falls short of his goal. Notice that this criterion may be meaningful also in formally complete markets if the initial capital of the investor is insufficient to achieve a perfect hedge. Notice also that the smallest V_0 for which, for given S_0, the minimal shortfall risk is zero, corresponds to the initial capital required for superhedging. Basic studies related to this criterion can be found in [18, 15, 32, 48].

Instead of the asymmetric criterion in (3.4) one may consider a symmetric criterion, namely.

C) Quadratic loss criterion. Given S_0, V_0, determine a self financing π such that

$$E_{S_0,V_0}\left\{(H_T - V_T^\pi)^2\right\} \quad \to \quad \min\,. \tag{3.5}$$

Compared to (3.4), criterion (3.5) may appear less convincing from an economic point of view, since it penalizes symmetrically losses as well as profits. On the other hand, with this symmetric penalization it makes sense to generalize the problem by minimizing also over V_0. Mathematically, criterion (3.5) is simpler to handle than (3.4) and may lead to computable solutions where (3.4) does not (see also Section 5.4 below). In the latter case one may determine a solution according to (3.5) and by simulations evaluate its performance with respect to other criteria as e.g. the asymmetric criterion of shortfall. If this performance turns out not to be satisfactory, then one may adjust the strategy heuristically to improve performance. Some basic studies related to this criterion may be found in [33, 34, 25, 75, 73].

D) Tracking a benchmark. The problem of hedging a liability corresponds to a finite horizon problem with only a terminal cost. A generalization that includes also a running cost is the problem of *tracking a benchmark* that, recalling the dynamics (2.3), (2.13) and (2.8) for the asset prices, benchmark and portfolio value, can be formulated as follows

$$\begin{cases} dS_t = (diag\, S_t)\, A_t dt + (diag\, S_t)\, \Sigma_t dw_t \\ dI_t = I_t\left[\alpha_t dt + \sigma_t dw_t + \rho_t dv_t\right] \\ dV_t^\pi = V_t^\pi\left[r_t dt + \pi_t(A_t - r_t\underline{1})dt + \pi_t \Sigma_t dw_t\right] - c_t dt \\ \max\limits_\pi E_{S_0,I_0,V_0}\left\{-\displaystyle\int_{T-\delta}^T \mathcal{L}(I_t, V_t^\pi)dt + \mathcal{G}(S_T, I_T, V_T^\pi)\right\} \end{cases} \tag{3.6}$$

where $0 < \delta \leq T$ and $\mathcal{L}(\cdot)$ and $\mathcal{G}(\cdot)$ may take the following form

$$\begin{aligned} \mathcal{L}(I_t, V_t^\pi) &= (g(I_t) - \eta - V_t^\pi)^+ \\ \mathcal{G}(S_T, I_T, V_T^\pi) &= V_T^\pi - H(I_T, S_T) \end{aligned} \tag{3.7}$$

for some functions $g(\cdot)$ and $H(\cdot)$ and for an $\eta > 0$. This implies that the investor tries to invest in a way that, at maturity T, the portfolio value exceeds as much as possible a function of the assets and the benchmark, without letting V_t^π fall much below a function of the benchmark in a left neighborhood $(T-\delta, T]$ of the maturity. It is clear that many variants of this example (3.6), (3.7) may be considered and for this we refer to [9].

Remarks. The basic problem setup, described in this subsection in different possible forms, can be generalized in various ways that correspond to the various possible generalizations of a standard stochastic control problem. Below we mention some of them.

- The dynamics of S_t, and thus also of I_t and V_t, may include a *discontinuous part* as in (2.12).
- The coefficients in the various models may be functions of a hidden Markov process as in (2.11) and, furthermore, depend on unknown parameters and this corresponds to *stochastic control under partial information* and *stochastic adaptive control*.
- Transaction costs of various kinds may be included, such as proportional transaction costs like in (2.10). With models of the type (2.10) one obtains *singular stochastic control problems* (see e.g. [78, 1]) as well as *impulse control* if one includes also fixed costs (see [50]). Besides transaction costs, one may also introduce various restrictions on the strategies themselves as well as on the value functions.
- *Impulse control problems* do not only arise with fixed transaction costs, but also under other circumstances such as interventions by a central bank to control the exchange rate (see e.g. [41, 11] and references therein). For a survey on impulse control problems in finance see [51].
- The basic hedging problem involves the price process S_t with initial condition S_0 as well as the value process V_t that, at maturity T, should possibly match the claim. This latter requirement imposes a form of terminal condition on V_t that depends on the evolution of S_t. This setup is typical for an application of techniques based on *forward-backward stochastic differential equations* (see e.g. [14, 28]).
- With a given maturity T, the hedging problem is naturally a finite horizon problem. Taking however into account that some of the assets may be subject to default, the horizon T could be replaced by a stopping time τ. This stopping time may also become

a control variable for the hedging problem of American-type options. The horizon may finally become also infinite for problems of *life-time consumption*, but also in portfolio optimization problems with the objective of maximizing the growth rate (see e.g. [1]). By maximizing the risk sensitized growth rate one ends up with problems of *risk sensitive asset management* (see e.g. [6, 31, 61]) that corresponds to an application of risk sensitive control.

4. Solution approaches. The control/optimization problems described in the previous section are all problems of the form of an optimal stochastic control problem or one of its variants. A natural solution approach is therefore based on the methodology of **dynamic programming (DP)** both for continuous as well as discrete time problems. As DP is a well known methodology, we shall not dwell on it here, but mention an alternative methodology, namely the "martingale method", that may occasionally be more convenient than DP. To introduce the intuition of this alternative methodology, in the next subsection we shall recall the standard approach to determine a self financing hedging strategy in a complete market.

4.1. Hedging in a complete market. Consider a market with N risky assets, for which the prices satisfy (2.3) with w_t an $N-$Wiener (i.e. we assume $M = N$). Let Σ_t be nonsingular for each t. In this situation the market is complete in the sense that, given any liability/claim, there exists an initial capital and a self financing strategy leading to a perfect hedge. Furthermore, there exists a unique equivalent *martingale measure (EMM)*, namely a measure $Q \sim P$ such that, under Q, the discounted prices $\tilde{S}_t^i = B_t^{-1} S_t^i$ are (local) martingales. It means that, under Q, the prices S_t have to satisfy, instead of (2.3),

$$dS_t = (diag\, S_t)\, r_t \underline{1} dt + (diag\, S_t)\, \Sigma_t dw_t^Q \tag{4.1}$$

where w_t^Q is a $(Q, \mathcal{F}_t)-$Wiener of dimension N. It follows in fact from (4.1) and Ito's rule that

$$d\tilde{S}_t = \left(diag\, \tilde{S}_t\right) \Sigma_t dw_t^Q. \tag{4.2}$$

Comparing (2.3) with (4.1) we have that

$$dw_t^Q = dw_t + \theta_t dt \quad \text{where} \quad \theta_t := \Sigma_t^{-1}(A_t - r_t \underline{1}) \tag{4.3}$$

i.e. the measure Q is obtained from P by a Girsanov-type measure transformation involving a translation of the Wiener process w_t by the process θ_t in (4.3) and so the Radon-Nikodym derivative $Z_t := (dQ/dP)_{|\mathcal{F}_t}$ satisfies (under P)

$$dZ_t = -Z_t \theta_t dw_t \tag{4.4}$$

From (4.2) and (2.9) it follows, furthermore, that also the discounted value $\tilde{V}_t$ of a self financing portfolio is a $(Q, \mathcal{F}_t)-$(local) martingale; in fact

$$d\tilde{V}_t = \phi_t d\tilde{S}_t = \phi_t (diag\, \tilde{S}_t) \Sigma_t dw_t^Q \tag{4.5}$$

which in integral form becomes

$$\tilde{V}_t = \tilde{V}_0 + \int_0^t \phi_s (diag\, \tilde{S}_s) \Sigma_s dw_s^Q . \tag{4.6}$$

Given a liability of the form $H_T = H(S_T)$, the hedging problem consists in finding $V_0 = \tilde{V}_0$ and ϕ_t such that $\tilde{V}_T = \tilde{H}_T$ a.s. where $\tilde{H}_T = B_T^{-1} H_T$. For this purpose consider, besides $\tilde{V}_t$, also the $(Q, \mathcal{F}_t)-$martingale

$$\tilde{M}_t := E^Q\{\tilde{H}_T | \mathcal{F}_t\} = E^Q\{\tilde{H}_T | \tilde{S}_t\} := F(t, \tilde{S}_t) \tag{4.7}$$

where we use E^Q to denote expectation with respect to Q (we keep denoting simply by E the expectation with respect to P) and where the second equality follows from $\tilde{H}_T = B_T^{-1} H(S_T)$ and the fact that S_t is Markov, while the last equality defines implicitly the function $F(t, s)$. If this latter function is sufficiently smooth, by Ito's rule we have, under Q,

$$\begin{aligned} d\tilde{M}_t &= dF(t, \tilde{S}_t) \\ &= \left[F_t(\cdot) + \frac{1}{2} tr\{(diag\, \tilde{S}_t) \Sigma_t \Sigma_t' (diag\, \tilde{S}_t)\} F_{ss}(\cdot) \right] dt \\ &\quad + F_s(\cdot)(diag\, \tilde{S}_t) \Sigma_t dw_t^Q . \end{aligned} \tag{4.8}$$

Since $\tilde{M}_t$ is by definition a martingale, the finite variation terms in (4.8) have to vanish implying that $F(t, s)$ has to satisfy the following PDE in $(t, s) \in \mathbb{R}^+ \times \mathbb{R}^N$

$$\begin{cases} F_t(t, s) + \dfrac{1}{2} tr\{(diag\, s) \Sigma_t \Sigma_t' (diag\, s)\} F_{ss}(t, s) = 0 \\ F(T, s) = \tilde{H}(s). \end{cases} \tag{4.9}$$

In integrated form, (4.8) then becomes

$$\tilde{M}_t = \tilde{M}_0 + \int_0^t F_s(t, \tilde{S}_t)(diag\, \tilde{S}_t) \Sigma_t dw_t^Q . \tag{4.10}$$

It follows that, in order to determine V_0 and ϕ_t such that $\tilde{V}_T^\phi = \tilde{H}_T$, it suffices to choose them so that $\tilde{V}_t = \tilde{M}_t$ and this is achieved by choosing

$$\begin{cases} \tilde{V}_0 = \tilde{M}_0 = E^Q\{\tilde{H}_T\} \\ \phi_t = F_s(t, \tilde{S}_t). \end{cases} \tag{4.11}$$

Notice that (4.10) may be viewed as an explicit form of the martingale representation for the $(Q, \tilde{F}_t)-$ martingale $\tilde{M}_t$ and this martingale representation leads via (4.11) to the solution of the hedging problem.

The idea of using martingale representation to solve the hedging problem has led to the "martingale approach", alternative to DP, to solve general stochastic control problems. It consists of two steps:
1) Determine the optimal value of the cost functional.
2) Determine the control/strategy that achieves this optimal value.
While for step 2) one may proceed as in the hedging problem via a *martingale representation*, for step 1) a possible methodology is mentioned in the next subsection.

4.2. The martingale method/convex duality to determine the optimal value. We recall here this method in its classical form and in the context of maximizing utility from terminal wealth (class A) in Section 3.1), which is also the class of problems where the method is mostly applied. For a more comprehensive treatment of the subject we refer to [45, 49, 72].

Consider again the complete market model of Section 4.1, that admits the unique martingale measure Q with Radon-Nikodym derivative $Z_t := (dQ/dP)_{|\mathcal{F}_t}$ satisfying (4.4) where θ_t is given in (4.3). Recall also from Section 4.1 (see (4.5),(4.6)) that, for any self financing strategy ϕ_t, the corresponding discounted value process $\tilde{V}_t^\phi$ is a $(Q, \mathcal{F}_t)-$martingale so that

$$\tilde{V}_0 (= V_0) = E^Q\{\tilde{V}_t\} = E\{Z_t B_t^{-1} V_t\}\,, \quad \forall t > 0. \tag{4.12}$$

Recall furthermore that, given an initial capital V_0, the problem of utility maximization from terminal wealth consists in finding a self financing strategy π_t (equivalently ϕ_t) that maximizes $E\{U(V_T^\pi)\}$ and so this problem can now be formulated as

$$\begin{cases} \max\limits_\pi E\{U(V_T^\pi)\} \\ E\{Z_T B_T^{-1} V_T^\pi\} = V_0 \end{cases} \tag{4.13}$$

where the second relation will be called **budget equation**. This formulation of the problem will be called **martingale formulation** and, in the form in which it is presented, it requires implicitly the existence of a unique martingale measure, i.e. a unique Radon-Nikodym derivative Z_t. Strictly speaking it applies thus only to complete markets (see however the Remark 4.1 at the end of this subsection as well as point ii) in the next subsection 4.3).

Next we recall two possible approaches to determine the optimal value.

4.2.1. Lagrange-multiplier method. Given (4.13) and dropping the superscript π, consider the Lagrangian

$$L(V_T, \lambda) = E\left\{U(V_T) - \lambda[Z_T B_T^{-1} V_T - V_0]\right\}. \tag{4.14}$$

Assuming that all operations can be performed, by differentiating with respect to V_T and λ one obtains the system

$$\begin{cases} E\left\{U'(V_T) - \lambda Z_T B_T^{-1}\right\} = 0 \\ V_0 - E\{Z_T B_T^{-1} V_T\} = 0. \end{cases} \tag{4.15}$$

Putting

$$I := (U')^{-1}, \tag{4.16}$$

the first equation leads one to choose

$$V_T = I(\lambda Z_T B_T^{-1}) \tag{4.17}$$

which, substituted into the second (budget equation), gives

$$V_0 = E\left\{Z_T B_T^{-1} I(\lambda Z_T B_T^{-1})\right\} := V(\lambda). \tag{4.18}$$

If this $V(\cdot)$ admits an inverse so that $\lambda = V^{-1}(V_0)$, then this λ, inserted into (4.17), gives the optimal terminal value as

$$V_T^* = I\left(V^{-1}(V_0) Z_T B_T^{-1}\right) \tag{4.19}$$

and so the optimal value in (4.13) is

$$V^* = E\{U(V_T^*)\} = E\left\{V^{-1}(V_0) Z_T B_T^{-1}\right\}. \tag{4.20}$$

In the particular case of $U(v) = \log v$ one has $V(\lambda) = \lambda^{-1}$ and so $V_T^* = V_0 Z_T^{-1} B_T$.

4.2.2. Convex duality. Consider the Legendre-Fenchel-type transformation (*convex dual*)

$$\tilde{U}(y) := \max_x [U(x) - xy] = U(I(y)) - yI(y) \tag{4.21}$$

from which

$$U(I(y)) - yI(y) \geq U(x) - xy\ , \quad \forall x. \tag{4.22}$$

Coming back to problem (4.13), consider any V_T satisfying the budget equation. Furthermore, put

$$V_T^* = I(\lambda Z_T B_T^{-1}) \tag{4.23}$$

with λ chosen so that also V_T^* satisfies the budget equation, namely

$$E\left\{Z_T B_T^{-1}\, I(\lambda Z_T B_T^{-1})\right\} = V_0. \tag{4.24}$$

Notice now that the left hand side in (4.24) coincides with $V(\lambda)$ in (4.18) so that, in the case of a logarithmic utility, $\lambda = V_0^{-1}$.

Putting $x = V_T$ and $y = \lambda Z_T B_T^{-1}$, (4.22) with (4.23) leads to

$$U(V_T^*) - \lambda Z_T B_T^{-1} V_T^* \geq U(V_T) - \lambda Z_T B_T^{-1} V_T. \tag{4.25}$$

Taking expectations on both sides of (4.25), thereby recalling (4.24) and the fact that both V_T^* and V_T satisfy the budget equation, one obtains

$$E\{U(V_T^*)\} - \lambda V_0 \geq E\{U(V_T)\} - \lambda V_0 \tag{4.26}$$

showing that $E\{U(V_T^*)\} \geq E\{U(V_T)\}$ for any V_T satisfying the budget equation. The value V_T^* in (4.23) is thus indeed the optimal terminal value.

REMARK 4.1. The martingale method/convex duality has been recalled here in its simplest form that requires implicitly the completeness of the market. The method has however been extended also to the case of intertemporal consumption, transactions costs and incomplete markets as well as to the case when there are constraints on the strategy and the value function. For this we refer e.g. to [45, 49] and [72] as general references and to [16, 17, 53, 19, 18, 15, 22, 8, 71, 47] in particular.

4.3. Remarks concerning the DP and the martingale/convex duality approaches. Although there are some similarities between the two approaches, they are rather different both in their basic setup as well as in their use.

i) In continuous time, the DP method is based on the Hamilton-Jacobi-Bellman (HJB) equation: first one determines the optimal control as a function of the (yet unknown) optimal value; substituting this control back into the HJB equation one obtains a nonlinear PDE, whose solution then leads to the optimal value. In the martingale method one proceeds in the inverse direction and the two problems of determining the optimal value and the optimal control are less intertwined: first one determines the optimal value without reference to the control and then the optimal strategy is determined as a strategy that leads to this optimal value.

ii) In its essence, DP is a fully dynamic approach and its important aspect (both in continuous and discrete time) is that, provided that the state process is Markovian and the cost structure is "additive" over time, it reduces optimization over time to a parameter optimization. This makes the problem of determining the optimal dynamic control easier; the more difficult problem in continuous time is the solution of the nonlinear PDE to obtain the optimal value (in discrete time one faces the well-known problem of the *curse of dimensionality*). With the martingale method the determination of the optimal value may be easier but, besides the fact that, contrary to DP, it does not require Markovianity, the method has in general a narrower field of applicability: it is tailored to more specific situations, in particular when the goal is utility maximization/loss minimization and there exists a martingale measure under which the value of a self financing portfolio becomes a martingale. In the classical setup recalled in Section 4.2 one needs furthermore the existence of the inverse of the derivative of the utility function as well as of the function $V(\lambda)$ in (4.18). (See however the generalizations mentioned in Remark 4.1).

iii) Explicit solutions are generally not easy to obtain by either of the methods. Rather powerful approximation methods have therefore been developed for DP and these methods can also be adapted to obtain stability/robustness results. For the martingale approach this seems not to have been done so far. On one hand it may be more difficult due to the martingale requirement, on the other hand the neater separation of the problems of determining the optimal value and the optimal control may be an advantage.

iv) A situation where, at least so far, DP seems to be better suited than the martingale method, is that of partial/incomplete information. With DP one solves to so-called *separated problem* where the unknown/unobserved quantities are replaced by their distributions that, using a stochastic filtering algorithm, are continuously updated on the basis of the successive information coming from the observations (in this context see also Section 5.2 below). This is basically possible due to the inherently dynamic nature of DP (the required Markov property for the information process is generally not a problem), while the martingale method is more static in nature. In Section 6 of [18] and in [46] (see also the previous study in [54]) one can find an example of application of the martingale method to a case of incomplete information. It concerns however the very particular setting when the uncertainty is in the stock appreciation rates that are supposed to be unknown constants and treated, from the Bayesian point of view, as random variables with a given prior distribution. In such a setting, the stochastic dynamic filtering that one would apply in a more general setup collapses to the static Bayesian updating and so the (static) duality approach becomes applicable (for a previous study in a more specifically Bayesian context see [10]). An extension to non constant, mean-reverting diffusion drift processes can be found in [83], where the author uses a technique based on the inverse Laplace transform to obtain explicit results. In this context we mention also [77] where, for a specific linear-Gaussian factor model, a utility maximization problem is approached by the method of convex duality. Finally, for the case of a risk-sensitive criterion under partial information, we refer to [61].

In the next section we shall discuss various issues of partial/incomplete information and this mainly in the context of the hedging problem.

5. Incomplete information, model uncertainty. For the solution of any applied problem, in particular a financial problem, one needs first to build a model. The solution may be rather sensitive to the model. On the other hand, one has often only incomplete/partial knowledge of the true model that, furthermore, may be considerably complex. This leads one to deal with the issue of model uncertainty that in financial applications corresponds to what is called **model risk**. In Subsections 5.1-5.3 below we shall mention three possible approaches to the hedging problem under model uncertainty and conclude the section in 5.4 by mentioning how for the case of a quadratic loss function (case C) in Section 3.2) the search for a solution under incomplete information simplifies considerably.

5.1. Min-max approach. This rather natural approach corresponds to the homonymous approach in decision theory and is a conservative approach in the sense that it protects the decision maker against the worst case scenario.

A general study of this approach can be found in [48] (previous studies appear in [18] and [15]) where, for the problem of shortfall risk minimization (criterion B) in Section 3.2), a family $\mathcal{P}$ of possible *real world probability measures* is considered and as criterion one uses

$$\inf_{\pi} \sup_{P \in \mathcal{P}} E^{P}_{S_0,V_0} \left\{ \mathcal{L} \left[(H_T - V_T^{\pi})^+ \right] \right\} \tag{5.1}$$

that reflects the aim of an agent to minimize the worst case risk. In (5.1) we now write the superscript P in E^P since we consider various measures $P \in \mathcal{P}$. The min-max quantity in (5.1) may also be considered as upper value of a fictitious stochastic game between the market and the agent and so it is of interest to find out whether this game has a value, i.e. whether the upper min-max value in (5.1) coincides with the lower max-min value

$$\sup_{P \in \mathcal{P}} \inf_{\pi} E^{P}_{S_0,V_0} \left\{ \mathcal{L} \left[(H_T - V_T^{\pi})^+ \right] \right\} . \tag{5.2}$$

The equality of (5.1) with (5.2) has been shown to hold true in full generality in [48], thus generalizing previous such results in [18] and[15]. In fact, in [18] the result is shown only for the case when the risk neutral measure is contained in $\mathcal{P}$ and in [15] when all the martingale measures are included in $\mathcal{P}$, i.e. when the uncertainty about the real world probability measure is large enough.

The idea of treating market uncertainty as a min-max problem is also picked up in [80] where the authors consider the model risk control problem as a two-player (trader vs market) zero-sum stochastic differential game problem. By allowing for uncertainty in the drift and the volatility of a standard log-normal model, it is shown in [80] that the min-max value of the game is the unique viscosity solution of a Hamilton-Jacobi-Bellman-Isaacs equation and that it satisfies a DP principle. This characterization allows the authors in [80] to develop also numerical methods.

The min-max approach has features that are similar to the superhedging criterion (criterion A) in Section 3.2); in particular, in order to achieve a satisfactorily small value of the shortfall risk, it requires a large initial capital. Furthermore, this approach does not allow to incorporate successive information on the model that becomes available to an agent by observing the prices on the market. In the next subsection we shall therefore discuss an approach that takes this feature into account.

5.2. Adaptive approaches (Stochastic control under partial information). An alternative way to deal with model uncertainty is to consider parameterized families of models and to update successively the knowledge about the parameter values on the basis of observed quantities such as prices of traded assets. The "parameters" may more generally be hidden processes such as the process Z_t in the Markov modulated model (2.11), or combinations of unknown parameters and hidden processes. Taking the Bayesian point of view, whereby an unknown quantity is considered a realization of a random variable, updating the knowledge of a parameter value corresponds to updating its distribution. For a combination of unknown parameters and hidden processes this leads to a combined filtering and parameter estimation problem, where the observations are the prices S_t. The unobserved/unknown quantities are then replaced by their current updated distributions that become so part of the state (which is finite-dimensional if the filter admits a finite dimensional representation) of the separated problem, mentioned under point iv) in Section 4.3, and to which DP can then be applied.

For the rest of this section we shall discuss the above setup in the context of discrete time models referring mainly to [69] and [68] (see however also [48]). To this effect recall that, discretizing in time according to (2.15) a Markov modulated model of the form (2.11), one obtains a sequence of random variables ξ_n^i that are lognormal conditional on Z_t. If Z_t is an unobserved/hidden process and model (2.11) contains additional unknown parameters, then the ξ_n^i in (2.15) are lognormal with unknown mean and variance. If one furthermore discretizes in space (*quantization*) the lognormal variables ξ_n^i, then (see always the discussion after (2.15)) one ends up with a multinomial price evolution model, where the multinomial probabilities are not fully known. Such models are thus rather general models in discrete time for which, in the context of shortfall risk minimization, an adaptive approach has been studied in [69] and [68] that we briefly recall now.

Considering for simplicity a single risky asset, i.e. $S_{n+1} = S_n\xi_n$, assume that ξ_n is $K-$nomial and denote by $\xi_n^{(k)}$ and $p^{(k)}$ its $k-$th ($k = 1,\cdots,K$) possible value and corresponding probability respectively. For the evolution of the portfolio process V_t^ϕ, that corresponds to a self financing strategy ϕ_t, one may, in accordance with (2.16) and neglecting transaction costs, use the shorthand notation $V_{n+1}^\phi = V_{n+1}(V_n^\phi, S_n, \xi_n, \phi_{n+1})$. Denoting then by $J_n(S_n, V_n^\phi)$ the optimal cost-to-go in period n, the DP recursions become, in the case of full knowledge of the multinomial probabilities $\bar{p} = (p^{(1)},\cdots,p^{(K)})$,

$$
\begin{aligned}
&J_n(S_n, V_n^\phi) \\
&\quad = \inf_{\phi_{n+1}\in A(S_n,V_n^\phi)} \sum_{k=1}^K p^{(k)} J_{n+1}(S_n\xi_n^{(k)}, V_{n+1}(V_n^\phi, S_n, \xi_n^{(k)}, \phi_{n+1}))
\end{aligned}
\tag{5.3}
$$

where $A(S_n, V_n^\phi)$ is the set of admissible controls in period n that takes into account the requirement of self financing, as well as possible further restrictions.

Assume now that one has incomplete information represented by a subfiltration $\tilde{\mathcal{F}}_n \subset \mathcal{F}_n$; more specifically, assume that $\bar{p}$ is not known, but S_n and V_n are observed and ϕ_{n+1} is $\tilde{\mathcal{F}}_n-$adapted. It is then easily seen that the DP recursions (5.3) remain the same except that each $p^{(k)}$, $(k = 1, \cdots, K)$, has to be replaced in the generic period n by $E\{p^{(k)} | \tilde{\mathcal{F}}_n\}$. In other words, for this kind of problems a CE (certainty equivalence) property holds and this then allows the case of partial information to be handled at the same level of complexity as that of full information. In [68] it is shown that, implementing the above in the case of a binomial model where the initial capital is not sufficient for a perfect hedge, one arrives at an interesting intuitive generalization of the classical CRR strategy (see [13]).

We close this subsection by pointing out that, if instead of a shortfall risk criterion we had considered the superhedging criterion, then the optimal strategy would be the same for all values of the multinomial probabilities $\bar{p} = (p^{(1)}, \cdots, p^{(K)})$ so that no adaptive approach would be necessary. In fact, since the superhedging criterion is an a.s. criterion, all models that induce equivalent probability measure lead to the same solution.

5.3. Robust approaches. An alternative way to deal with model uncertainty is to investigate the sensitivity of the solution with respect to the model or, in other words, the robustness of the solution. One may in fact ask how reliable is the solution obtained for a hypothetical model, when applied to the real model.

To make this more formal, suppose that one has a generic finite horizon stochastic control problem with running and terminal cost that is defined on a filtered probability space $(\Omega, \mathcal{F}, \mathcal{F}_t, P)$, and where the (controlled) state is denoted by X_t and the control by π_t. Suppose that the dynamics are driven by an $(\mathcal{F}_t, P)-$Wiener process and are affected also by a (hidden) Markov process Z_t, i.e.

$$(5.4) \qquad \begin{cases} dX_t = A_t(X_t, Z_t, \pi_t)dt + \Sigma_t(X_t, Z_t, \pi_t)dw_t \\ J^P(X_0, \pi^*) = \inf\limits_{\pi} J^P(X_0, \pi) \\ \qquad\qquad\quad = \inf\limits_{\pi} E^P_{X_0}\Big\{ \int_0^T c(X_t, \pi_t)dt + C(X_T) \Big\}. \end{cases}$$

This problem may, in particular, be the utility maximization problem (3.1) (with $X_t := [S_t, V_t]'$), or the corresponding hedging problem with criterion (3.4) or (3.5), or the benchmark tracking problem (3.6) (with $X_t := [S_t, I_t, V_t]'$), but where S_t satisfies the Markov modulated model (2.11) instead of (2.3). We shall refer to (5.4) as the *real world problem* that corresponds to the *real world measure* P. The control π^* denotes the optimal control whenever it exists (what matters is the optimal minimal value).

Assume now that one does not know exactly the real world measure P and that one solves the same problem, but for a hypothetical measure Q, that is absolutely continuous (equivalent) with respect to P. We want to point out here that, in diffusion-type models, the change to an absolutely continuous measure corresponds to a translation of the Wiener process, i.e. to a change in the drift. This is in fact the almost exclusive situation considered in the mathematical finance literature that deals with incomplete information (see point iv) in Section 5.3 and, in particular, [54, 63] and references therein; see however also [64] as well as [4, 55] and [80] for uncertainty in the volatility). In our more flexible model (5.4) one can however consider also absolutely continuous measures that do not necessarily correspond to a change in the drift. In fact, if Z_t is independent of w_t, then the probability space $(\Omega, \mathcal{F}, \mathcal{F}_t, P)$ may be viewed as a product space $(\Omega = \Omega^w \times \Omega^Z,\ P = P^w \times P^Z)$ and so one may consider an equivalent measure $Q = Q^w \times Q^Z$ where $Q^w \equiv P^w$ while $Q^Z \sim P^Z$, $(Q^Z \neq P^Z)$, where $\sim$ denotes equivalence between the measures. Assuming then a measure Q of this type, the Wiener process w_t remains the same under P and Q and only the distribution of Z_t changes. In this situation the hypothetical problem corresponding to (5.4) becomes

$$(5.5)\qquad \begin{cases} dX_t = A_t(X_t, Z_t, \pi_t)dt + \Sigma_t(X_t, Z_t, \pi_t)dw_t \\ J^Q(X_0, \pi^Q) = \inf_\pi E^Q_{X_0}\Big\{ \int_0^T c(X_t, \pi_t)dt + C(X_T)\Big\} \end{cases}$$

for which the optimal control is now denoted by π^Q (if such an optimal control does not exist, one may consider an $\epsilon-$optimal control instead).

The following question arises now naturally

Question: can one find a bound (uniform in X_0) on

$$(5.6)\qquad \Lambda_0^Q(X_0) := J^P(X_0, \pi^Q) - J^P(X_0, \pi^*) \geq 0$$

i.e. on the suboptimality of π^Q, when it is applied to the real problem, in terms of a measure of the difference between P and Q?

This question is dealt with in [30] in the context of discrete time control problems and here we shall briefly recall one of the results. Assuming X_0 fixed, by adding and subtracting it is easily seen that

$$(5.7)\qquad J^P(\pi^Q) - J^P(\pi^*) \leq 2 \sup_\pi \left| J^P(\pi) - J^Q(\pi) \right|$$

so that it suffices to have a bound on the right hand side in (5.7). Let $\delta > 0$ be such that

$$(5.8)\qquad \sup_{\omega \in \Omega} \left| 1 - \frac{dP(\omega)}{dQ(\omega)} \right| \leq \delta.$$

Then (see [30]) one obtains that

i) $\left|J^P(\pi) - J^Q(\pi)\right| \leq \delta \sup_\pi J^Q(\pi)$

ii) for $\delta < 1$ the bound in i) can be improved to become

$$\left|J^P(\pi) - J^Q(\pi)\right| \leq \frac{\delta(1+\delta)}{1-\delta} J^Q(\pi^Q).$$

The interesting aspect of ii) is that, when it applies, the solution of the hypothetical problem leads not only to a control π^Q, but also to a bound on its suboptimality.

These ideas have been applied in [29] to derive a "most robust" adaptive investment strategy for a discrete time binomial model, i.e. an adaptive strategy that minimizes $\Lambda_0^Q(X_0)$ in (5.6). This strategy turns out to be superior to the adaptive strategy in [68] in the sense that it induces much smaller excursions of the portfolio value before maturity.

5.4. The case of a quadratic loss function. It has been mentioned (see point C) in Section 3.2) that the quadratic loss has some mathematically attractive features and the purpose of this section is to mention some of these.

We shall in fact summarize here the main ideas from [76, 24, 36] (see also [63]). Assume we have a market with a non-risky and a risky asset, the latter being defined on a filtered probability space $(\Omega, \mathcal{F}, \mathcal{F}_t, Q)$ with Q a martingale measure. The volatility is furthermore supposed to be driven, as in (2.11), by a (hidden) Markov process Z_t so that, under Q, we may write

$$dS_t = S_t\sqrt{v(t, S_t, Z_t)}dw_t. \tag{5.9}$$

If Z_t is adapted to a larger filtration than $\mathcal{F}_t = \mathcal{F}_t^w$, say $\mathcal{G}_t = \mathcal{F}_t^{w,Z}$, then even with full information of Z_t, the market is incomplete and perfect hedging with a self financing strategy is no longer possible. If one chooses as hedging criterion the quadratic loss (criterion C) in Section 3.2), then a strategy that minimizes the corresponding risk can be found on the basis of a Kunita-Watanabe decomposition (see e.g. [34], see also [73], and [66] for a corresponding discrete time version), that is a generalization of the well-known martingale representation theorem and, as the latter, can in many cases be obtained by an application of the Ito formula. Let then $\phi_t^*(S_t, Z_t)$ be the optimal strategy under full information of S_t and Z_t, which describes the number of shares of the risky asset to be kept in the portfolio at time t. Thanks to the quadratic nature of the criterion, the following can be shown (see the references at the beginning of this subsection; see also [66] for the discrete time case): the strategy minimizing the risk under the partial information corresponding to the subfiltration $\mathcal{F}_t^S \subset \mathcal{F}_t^{w,Z}$ (resulting from the observation only of the prices S_t of the risky asset) is obtained by "projecting" $\phi_t^*(S_t, Z_t)$ onto this subfiltration $\mathcal{F}_t^S$. More precisely, for the optimal strategy ϕ_t^* under partial observation one has (see [36] for the case of a volatility as in (5.9))

$$\phi_t^* = \frac{E\left\{v(t,S_t,Z_t)\phi_t^*(S_t,Z_t)\mid \mathcal{F}_t^S\right\}}{E\left\{v(t,S_t,Z_t)\mid \mathcal{F}_t^S\right\}} \tag{5.10}$$

which is a quantity that can be determined on the basis of the filter distribution $p(Z_t|\,\mathcal{F}_t^S)$.

6. Additional topics. We mention here two topics of recent investigation that may be of interest in the context of this survey. These two topics are by no means exhaustive, but reflect some of the specific interests of the author.

6.1. A case of almost sure optimality. As mentioned in the remarks to Section 3.2, for the problem of tracking a benchmark (Problem D) of Section 3.2) one may also consider an infinite horizon. The investment manager may in fact have as his goal the tracking of a given benchmark over a long period. Referring to model (3.6) one may then consider a quantity of the form

$$J_T(\pi) := \int_0^T c(V_t^\pi, I_t)dt \tag{6.1}$$

where $c(V_t^\pi, I_t)$ is some measure of discrepancy between the portfolio value V_t^π and the index I_t in period t. As criterion one may use the classical average cost per unit time

$$\gamma^\pi := \limsup_{T\to\infty} \frac{1}{T} E\{J_T(\pi)\} \tag{6.2}$$

or an almost sure criterion as follows (see [20]). Let P^π be the measure induced on the space of the trajectories of the state variable by a given strategy (think of the trajectories X_t in (5.4) where $X_t = [S_t, I_t, V_t^\pi]'$ with the components satisfying (3.6)). Furthermore, let $g(T)$ be a non-increasing function on $[0,+\infty)$ such that $\lim_{T\to\infty} g(T) = 0$. We shall say that a strategy π^* is $g-$a.s. optimal, if

$$\lim_{T\to\infty} g(T)\,[J_T(\pi^*) - J_T(\pi)]^+ = 0\,, \quad P^{\pi^*,\pi} - \text{a.s.}, \quad \forall\pi \tag{6.3}$$

where $P^{\pi^*,\pi}$ is any measure on the pairs of trajectories that has as marginals the measures P^{π^*} and P^π induced by the strategies π^* and π respectively. In words, one tries to obtain a strategy that on an infinite horizon reduces the set of trajectories for which $J_T(\pi^*) > J_T(\pi)$ to a set of probability zero; in addition, the rate of growth of the random variable $[J_T(\pi^*) - J_T(\pi)]^+$ is controlled by the function $g(\cdot)$. A first approach to this problem has been developed in [81] and is currently being investigated further.

6.2. Factor models with unobserved factors (an application of stochastic filtering). Model (2.11) with Markov modulated coefficients is an instance of a factor model with the factors represented by the process Z_t. As mentioned, there may be various factors, each with an economic interpretation; some of them may be observable, some not. The prices on the market depend however on all the factors and so, when it comes to pricing and hedging, one has to explicitly describe the partial information that one has concerning the unobserved factors and this is where stochastic filtering comes into play. The main problem that arises in this context is that, for the pricing and hedging of derivatives, one has to work under a risk neutral martingale measure, while for the filtering as such one has to consider the real world probability measure.

This problem has been studied in the recent paper [65], where a method is presented that allows the derivative pricing to be performed under the real world probability measure. To this effect notice that a martingale measure is linked to a so-called *numeraire* (an asset or self financing portfolio with positive value) in the sense that the prices on the market, discounted with respect to a given numeraire, are martingales with respect to the corresponding martingale measure. Among the possible numeraire portfolios a special role is plaid by the so-called *growth-optimal portfolio* (GOP), which is the self financing portfolio that maximizes expected logarithmic utility. It turns out that one can obtain a fair pricing system so that the corresponding prices, expressed in units of GOP, are indeed local martingales under the given real world probability measure.

The factors, considered in [65], are represented as $z_t = (x_t, y_t)$, where x_t is the subvector of unobservable factors, y_t that of the observable ones and these factors are supposed to form a Markovian jump-diffusion model. Given an underlying filtered probability space $(\Omega, \mathcal{F}, \mathcal{F}_t, P)$, consider then an observable subfiltration $\tilde{\mathcal{F}}_t \subset \mathcal{F}_t$. Denoting by $p(x_t | \tilde{\mathcal{F}}_t)$ the filter distribution of the unobservable factors, assume that the filter admits a finite-dimensional realization, i.e.

$$p(x_t | \tilde{\mathcal{F}}_t) = p(x_t; \xi_t) \tag{6.4}$$

where ξ_t is a finite-dimensional vector. The original factors $z_t = (x_t, y_t)$ can then be replaced by the observable factors $\tilde{z}_t = (\xi_t, y_t)$ and, under suitable assumptions, also $\tilde{z}_t$ turns out to have Markovian dynamics. It is furthermore assumed that in the market there is a certain number of primary assets (currencies or shares), which are supposed to be observable and to securitize the randomness in the market that comes from the driving Wiener and jump processes (the market is *complete*). With these assets one can form self financing portfolios, among them the GOP. The dependence of the asset dynamics on the factors comes from considering the value of the GOP as a function of the factors.

Finally, let $U(T, y_T)$ be a contingent claim at a given maturity T. It is not restrictive to let it depend only on the observable factors, since its value has to become known at T. Recall that prices, expressed in units of GOP, are local martingales under P. By assuming that $U(T, y_T)$ is integrable, in [65] the fair price of this claim is then defined as the $(P, \tilde{\mathcal{F}}_t)-$martingale

$$u(t, \tilde{z}_t) := E\left\{U(T, y_T) \,|\, \tilde{\mathcal{F}}_t\right\}. \tag{6.5}$$

Notice that, defining also the $(P, \mathcal{F}_t)-$martingale

$$u(t, z_t) := E\left\{U(T, y_T) \,|\, \mathcal{F}_t\right\} \tag{6.6}$$

namely the fair price under the hypothetical full information of all the factors, one has

$$u(t, \tilde{z}_t) = E\left\{u(t, z_t) \,|\, \tilde{\mathcal{F}}_t\right\} \tag{6.7}$$

which can be computed on the basis of the filter distribution $p(x_t | \tilde{\mathcal{F}}_t)$ and which also emphasizes the role of filtering for pricing under incomplete information.

In [65] the authors study also the hedging problem in the above context and it turns out that, given the assumption that the primary assets securitize the market, the hedging strategy becomes a function of the observable factors $\tilde{z}_t$. The filtered market model remains thus complete despite of the fact that the original model involves unobserved factors.

REFERENCES

[1] M. Akian, A. Sulem and M. Taksar, *Dynamic optimization of a long term growth rate for a mixed portfolio with transaction costs*, Mathematical Finance 11 (2001), pp. 153–188.

[2] J.P. Ansel and C. Srticker, *Couverture des actifs contingents et prix maximum*, Annales Institut Henri Poincaré 30 (1994), pp. 303–315.

[3] S. Asmussen, B. Højgaard and M. Taksar, *Optimal risk and dividend distribution policies. Example of excess-of-loss reinsurance for an insurance corporation*, Finance and Stochastics 4 (2000), pp. 299–342.

[4] M. Avellaneda, A. Levy and A. Paras, *Pricing and Hedging Derivative Securities in Markets With Uncertain Volatilities*, Applied Mathematical Finance 1 (1995), pp. 73–88.

[5] O.E. Barndorff-Nielsen, T. Mikosch and S. Resnick (eds), *Levy Processes: Theory and Applications*, Birkhäuser Verlag 2001.

[6] T.R. Bielecki and S.K. Pliska, *Risk sensitive dynamic asset management*, Applied Mathematics and Optimization 39 (1999), pp. 337–360.

[7] F. Black and M. Scholes, *The pricing of options and corporate liabilities*, Journal of Political Economics 81 (1973), pp. 637–659.

[8] B. Bouchard, *Utility maximization on the real line under proportional transaction costs*, Finance and Stochastics 6 (2002), to appear.

[9] S. Browne, *Beating a moving target: Optimal portfolio strategies for outperforming a stochastic benchmark*, Finance and Stochastics 3 (1999), pp. 275–294.

[10] S. Browne and W. Whitt, *Portfolio choice and the Bayesian Kelly criterion*, Advances in Applied Probability 28 (1996), pp. 1145–1176.

[11] A. Cadenillas and F. Zapatero, *Classical and impulse stochastic control of the exchange rate using interest rates and reserves*, Mathematical Finance 10 (2000), pp. 141–156.

[12] T. Choulli, M. Taksar and X.Y. Zhou, *Excess-of-loss reinsurance for a company with debt liability and constraints on risk reduction*, Quantitative Finance 1 (2001), pp. 573–596.

[13] J.C. Cox, S.A. Ross and M. Rubinstein, *Option pricing: a simplified approach*, Journal of Financial Economics 7 (1979), pp. 229–263.

[14] J. Cvitanić, *Optimal trading under constraints*, In: Financial Mathematics (W.J. Runggaldier, ed.), Lecture Notes in Mathematics Vol. 1656, Springer Verlag, Berlin-Heidelberg 1997, pp. 123–190.

[15] J. Cvitanić, *Minimizing expected loss of hedging in incomplete and constrained markets*, SIAM Journal on Control and Optimization 38 (2000), pp. 1050–1066.

[16] J. Cvitanić and I. Karatzas, *Convex duality in constrained portfolio optimization*, Annals of Applied Probability 2 (1992), pp. 767–818.

[17] J. Cvitanić and I. Karatzas, *Hedging and portfolio optimization under transaction costs: a martingale approach*, Mathematical Finance 6 (1996), pp. 133–165.

[18] J. Cvitanić and I. Karatzas, *On dynamic measures of risk*, Finance and Stochastics 4 (1999), pp. 451–482.

[19] J.Cvitanić, W. Schachermayer and H. Wang, *Utility maximization in incomplete markets with random endowment.* Preprint, Institute of Financial and Actuarial Mathematics, Vienna Univ. of Technology, 1999.

[20] P. Dai Pra, G.B. Di Masi and B. Trivellato, *Pathwise optimality in stochastic control*, SIAM Journal on Control & Optimization 39 (2000), pp. 1540–1557.

[21] M.H.A. Davis and A.R. Norman, *Portfolio selection with transaction costs*, Mathematics of Operations Research 16 (1990), pp. 676–713.

[22] G. Deelstra, H. Pham and N. Touzi, *Dual formulation of the utility maximization problem under transaction costs*, Annals of Applied Probability, 11 (2001), pp. 1353–1383.

[23] F. Delbaen and W. Schachermayer, *A general version of the fundamental theorem of asset pricing*, Mathematische Annalen 300 (1994), pp. 463–520.

[24] G.B. Di Masi, E. Platen and W.J. Runggaldier, *Hedging of options under discrete observation on assets with stochastic volatility.* In: Seminar on Stochastic Analysis, Random Fields and Applications (E. Bolthausen, M. Dozzi, and F. Russo, eds.). Progress in Probability, Vol. 36, Birkhäuser Verlag, 1995, pp. 359–364.

[25] D. Duffie and H.R. Richardson, *Mean variance hedging in continuous time*, Annals of Applied Probability 1 (1991), pp. 1–15.

[26] T.E. Duncan, Y.Z. Hu and B. Pasik-Duncan, *Stochastic calculus for fractional Brownian motion, I. Theory*, SIAM Journal on Control and Optimization 38 (2000), 582–612.

[27] N. El Karoui and M.C. Quenez, *Dynamic programming and pricing of contingent claims in an incomplete market*, SIAM Journal on Control and Optimization 33 (1995), pp. 29–66.

[28] N.El Karoui and M.C. Quenez, *Non-linear pricing theory and backward stochastic differential equations.* In: Financial Mathematics (W.J. Runggaldier, ed.), Lecture Notes in Mathematics Vol. 1656, Springer Verlag, Berlin-Heidelberg 1997, pp. 191–246.

[29] G. Favero, *Shortfall risk minimization under model uncertainty in the binomial case: adaptive and robust approaches*, Mathematical Methods of Operations Research 53 (2001), pp. 493–503.

[30] G. Favero and W.J. Runggaldier, *A robustness result for stochastic control*, Systems and Control Letters. To appear.

[31] W.H. Fleming and S.J. Sheu, *Risk sensitive control and an optimal investment model*, Mathematical Finance 10 (2000), pp. 197–213.

[32] H. Föllmer and P. Leukert, *Efficient hedging: cost versus shortfall risk*, Finance and Stochastics 4 (2000), pp. 117–146.

[33] H. Föllmer and D. Sondermann, *Hedging of non-redundant contingemt claims.* In: Contributions to Mathematical Economics in Honor of Gérard Debreu (W. Hildenbrand and A. Mas-Colell, eds.), North-Holland, Amsterdam (1986), pp. 205–223.

[34] H. Föllmer and M. Schweizer, *Heding of contingent claims under incomplete information.* In: Applied Stochastic Analysis (M.H.A. Davis and R.J. Elliott, eds.), Gordon and Breach, London, New York (1991), pp. 389–414.

[35] R. Frey and W.J. Runggaldier, *A nonlinear filtering approach to volatility estimation with a view towards high frequency data*, International Journal of Theoretical and Applied Finance 4 (2001), pp. 199–210.

[36] R. Frey and W.J. Runggaldier, *Risk-minimizing hedging strategies under restricted information: the case of stochastic volatility models observed only at discrete random times*, Mathematical Methods of Operations Research 50 (1999), pp. 339–350.

[37] J. Gaier, P. Grandits and W. Schachermayer, *Asymptotic ruin probabilities and optimal investment.* Preprint, Institute of Financial and Actuarial Mathematics, Vienna Univ. of Technology, 2002.

[38] P. Guasoni, *Risk minimization under transaction costs*, Finance and Stochastics 6 (2002), pp. 91–113.

[39] C. Hipp and M. Plum, *Optimal investment for insurers*, Insurance: Mathematics and Economics 27 (2000), pp. 215–228.

[40] Y.Z. Hu, B. Øksendal and A. Sulem, *Optimal Portfolios in a Black and Scholes Market*, Preprint Univ. of Oslo 1999.

[41] M. Jeanblanc-Picqué, *Impulse control method and exchange rate*, Mathematical Finance 2 (1993), pp. 161–177.

[42] Yu. Kabanov and G. Last, *Hedging under transaction costs in currency markets: a continuous-time model*, Mathematical Finance 12 (2002), pp. 63–70.

[43] Yu. Kabanov and C. Stricker, *Hedging of contingent claims under transaction costs.* Preprint 2002.

[44] J. Kallsen, *Optimal portfolios for exponential Levy processes*, Mathematical Methods of Operations Research 51 (2000), 357–374.

[45] I. Karatzas and S.E. Shreve, *Methods of Mathematical Finance*, Springer-Verlag, New York, 1998.

[46] I. Karatzas and X. Zhao, *Bayesian adaptive portfolio optimization*, In: Handbook of Mathematical Finance, Cambridge Univ Press, pp. 632–670.

[47] I. Karatzas and G. Zitković, *Optimal consumption from investment and random endowment in incomplete semimartingale markets.* Preprint, Columbia University 2001.

[48] M. Kirch, *Efficient hedging in incomplete markets under model uncertainty.* PhD thesis, Humboldt University, Berlin, 2002.

[49] R. Korn, *Optimal Portfolios. Stochastic models for optimal investment and risk management in continuous time*, World Scientific Publishing Co. Pte. Ltd. 1997.

[50] R. Korn, *Portfolio optimization with strictly positive transaction costs and impulse control*, Finance and Stochastics 2 (1998), pp. 85–114.

[51] R. Korn, *Some applications of impulse control in mathematical finance*, Mathematical Methods of Operations Research 50 (1999), pp. 493–518.

[52] D. Kramkov, *Optional decomposition of supermartingales and hedging contingent claims in incomplete security markets*, Theory of Probability and Related Fields 105 (1996), pp. 459–479.

[53] D. Kramkov and W. Schachermayer, *The asymptotic elasticity of utility functions and optimal investment in incomplete markets*, Annals of Applied Probability 9 (1999), pp. 904–950.
[54] P. Lakner, *Utility maximization with partial information*, Stochastic Processes and Their Applications 56 (1995), pp. 247–273.
[55] T. Lyons, *Uncertain volatility and the riskfree synthesis of derivatives*, Applied Mathematical Finance 2 (1995), pp. 117–133.
[56] M.J.P. Magill and G.M. Constantinides, *Portfolio selection with transaction costs*, Journal of Economic Theory 13 (1976), pp. 245–263.
[57] B.B. Mandelbrot, *Fractals and Scaling in Finance*, Springer Verlag, 1997.
[58] A. Martin-Löf, *Lectures on the use of control theory in insurance*, Scandinavian Actuarial Journal 77 (1994), pp. 1–25.
[59] R. Merton, *Optimum consumption and portfolio rules in a continuous time model*, Journal of Economic Theory 3 (1971), pp. 373–413.
[60] R. Merton, *Option pricing when the underlying stock returns are discontinuous*, Journal of Financial Economics 5 (1976), pp. 125–144.
[61] H. Nagai and S. Peng, *Risk-sensitive dynamic portfolio optimization with partial information on an infinite time horizon*, Annals of Applied Probability, to appear.
[62] J. Paulsen and H. Gjessing, *Optimal choice of dividend barriers for a risk process with stochastic return of investment*, Insurance: Mathematics and Economics 20 (1997), pp. 215–223.
[63] H. Pham, *Mean-variance hedging for partially observed drift processes*, International Journal of Theoretical and Applied Finance 4 (2001), pp. 263–284.
[64] H. Pham and M.C. Quenez, *Optimal portfolio in a partially observed stochastic volatility model*, Annals of Applied Probability, to appear.
[65] E. Platen and W.J. Runggaldier, *A benchmark approach to filtering in finance.* Preprint 2002.
[66] W.J. Runggaldier, *The choice of the loss function for a decision problem under uncertainty as a compromise between representativity and tractability.* ES-ReDA International Seminar on Decision Analysis, Rome, June 18–19, 2001, pp. 141–148.
[67] W.J. Runggaldier, *Jump-diffusion Models*, In: Handbook of Heavy Tailed Distributions in Finance, North Holland Handbooks of Finance (Series Editor W.T. Ziemba), to appear.
[68] W.J. Runggaldier, B.Trivellato and T.Vargiolu, *A Bayesian adaptive control approach to risk management in a binomial model.* In: Seminar on Stochastic Analysis, Random Fields and Applications (R.C. Dalang, M. Dozzi, and F. Russo, eds.). Progress in Probability, Vol. 52, Birkhäuser Verlag, 2002, pp. 243–258.
[69] W.J. Runggaldier and A. Zaccaria, *A stochastic control approach to risk management under restricted information.* Mathematical Finance 10 (2000), pp. 277–288.
[70] W. Schachermayer, *The fundamental theorem of asset pricing under proportional transaction costs in finite discrete time*, Preprint 2001.
[71] W. Schachermayer, *Optimal investment in incomplete markets when wealth may become negative*, Annals of Applied Probability 11 (2001), pp. 694–734.
[72] W. Schachermayer, *Optimal investment in incomplete financial markets.* In: Mathematical Finance - Bachelier Congress 2000 (H. Geman, D. Madan, S. Pliska, and T. Vorst, eds.), Springer Verlag 2002, pp. 427–462.
[73] M. Schael, *On quadratic cost criteria for option hedging*, Mathematics of Operations Research 19 (1994), pp. 121–131.
[74] H.P. Schmidli, *On minimising the ruin probability by investment and reinsurance*, Annals of Applied Probability, to appear.
[75] M. Schweizer, *Mean variance hedging for general claims*, Annals of Applied Probability 2 (1992), pp. 171–179.

[76] M. Schweizer, *Risk minimizing hedging strategies under restricted information*, Mathematical Finance 4 (1994), pp. 327–342.

[77] J. Sekine, *Power-utility maximization for linear-Gaussian factor models under partial information*, Preprint, Osaka University, 2001.

[78] M.H. Soner and S. Shreve, *Optimal investment and consumption with transaction costs*, Annals of Applied Probability 4 (1994), pp. 609–692.

[79] H.M. Soner and N. Touzi, *Superreplication under Gamma constraints*, SIAM Journal on Control and Optimization 39 (2000), pp. 73–96.

[80] D. Talay and Zheng, *Worst case model risk management*, Finance and Stochastics, to appear.

[81] M. Tolotti, *Strategie di potafoglio ottimali per il "tracking" di un indice*, Laurea thesis, University of Padova 2002.

[82] T. Zariphopoulou, *Investment-consumption models with transaction fees and Markov-chain parameters*, SIAM Journal on Control and Optimization 30(1992), pp. 613–633.

[83] G. Zohar, *A generalized Cameron-Martin formula with applications to partially observed dynamic portfolio optimization*, Mathematical Finance 11 (2001), pp. 475–494.